

Surgical and Interventional Robotics

Core Concepts, Technology, and Design

BY PETER KAZANZIDES, GABOR FICHTINGER, GREGORY D. HAGER,
ALLISON M. OKAMURA, LOUIS L. WHITCOMB, AND RUSSELL H. TAYLOR

Two decades after the first reported robotic surgical procedure [1], surgical robots are just beginning to be widely used in the operating room or interventional suite. The da Vinci telerobotic system (Intuitive Surgical, Inc.), for example, has recently become more widely employed for minimally invasive surgery [2]. This article, the first in a three-part series, examines the core concepts underlying surgical and interventional robots, including the potential benefits and technical approaches, followed by a summary of the technical challenges in sensing, manipulation, user interfaces, and system design. The article concludes with a review of key design aspects, particularly in the areas of risk analysis and safety design. Note that medical care can be delivered in a surgical suite (operating room) or an interventional suite, but for convenience, we will henceforth use the term *surgical* to refer to both the surgical and interventional domains.

Core Concepts

This section describes some of the potential benefits of surgical robots, followed by an overview of the two technical paradigms, surgical computer-aided design and computer-aided manufacturing (CAD/CAM) and surgical assistance, which will be the subjects of the second and third articles in this series.

Potential Benefits

The development of surgical robots is motivated primarily by the desire to enhance the effectiveness of a procedure by coupling information to action in the operating room or interventional suite. This is in contrast to industrial robots, which were

developed primarily to automate dirty, dull, and dangerous tasks. There is an obvious reason for this dichotomy: medical care requires human judgment and reasoning to handle the variety and complexity of human anatomy and disease processes. Medical actions are chosen based on information from a number of sources, including patient-specific data (e.g., vital signs and images), general medical knowledge (e.g., atlases of human anatomy), and physician experience. Computer-assisted interventional systems can gather and present information to the physician in a more meaningful way and, via the use of robots, enable this information to influence the performance of an intervention, thereby potentially improving the consistency and quality of the clinical result. It is, therefore, not surprising that surgical robots were introduced in the 1980s, after the dawn of the information age, whereas the first industrial robot was used in 1961.

There are, however, cases where surgical robots share potential benefits with industrial robots and teleoperators. First, a robot can usually perform a task more accurately than a human; this provides the primary motivation for surgical CAD/CAM systems, which are described later in the “Surgical CAD/CAM” section. Second, industrial robots and teleoperators can work in areas that are not human friendly (e.g., toxic fumes, radioactivity, or low-oxygen environments) or not easily accessible to humans (e.g., inside pipes, the surface of a distant planet, or the sea floor). In the medical domain, inhospitable environments include radiation (e.g., X-rays) and inaccessible environments include space-constrained areas such as the inside of a patient or imaging system. This also motivates the development of surgical CAD/CAM systems and is one of the primary motivations for surgical assistant systems, described in the “Surgical Assistance” section.

Digital Object Identifier 10.1109/MRA.2008.926390

In contrast to industrial robots, surgical robots are rarely designed to replace a member of the surgical or interventional team. Rather, they are intended to augment the medical staff by imparting superhuman capabilities, such as high motion accuracy, or to enable interventions that would otherwise be physically impossible. Therefore, methods for effective human-robot cooperation are one of the unique and central aspects of medical robotics.

Technical Paradigms

In our research, we find it useful to categorize surgical robots as surgical CAD/CAM or surgical assistance systems, based on their primary mode of operation [3]. Note, however, that these categories are not mutually exclusive and some surgical robots may exhibit characteristics from both categories. The following sections briefly describe these categories, with representative examples.

Surgical CAD/CAM

The basic tenet of CAD/CAM is that the use of a computer to design a part creates a digital blueprint of the part, and so it is natural to use a computer-controlled system to manufacture it, i.e., to translate the digital blueprint into physical reality. In the medical domain, the planning that is often performed prior to, or during, an intervention corresponds to CAD, whereas the intervention represents CAM. To take the analogy further, postoperative assessment corresponds to total quality management (TQM). We refer to the closed-loop process of 1) constructing a patient-specific model and interventional plan; 2) registering the model and plan to the patient; 3) using technology to assist in carrying out the plan; and 4) assessing the result, as surgical CAD/CAM, again emphasizing the analogy between computer-integrated medicine and computer-integrated manufacturing (Figure 1).

The most well-known example of a surgical CAD/CAM system is ROBODOC (ROBODOC, a Curexo Technology Company; formerly Integrated Surgical Systems, Inc.) [4], [5]. ROBODOC was developed for total hip and total knee replacement surgeries (Figure 2). In these surgeries, the patient's joint is replaced by artificial prostheses: for hip surgery, one prosthesis is installed in the femur and another in the acetabulum (pelvis) to create a ball and socket joint; for knee surgery, one prosthesis is installed in the femur and the other in the tibia to create a sliding hinge joint. Research on ROBODOC began in the mid-1980s as a joint project between IBM and the University of California, Davis. At that time, the conventional technique for hip and knee replacement surgery consisted of two-dimensional (2-D) planning

(using X-rays) and manual methods (handheld reamers and broaches) for preparing the bone. The motivation for introducing a robot was to improve the accuracy of this procedure—both the placement accuracy (to put the prostheses in the correct places) and the dimensional accuracy (to get a good fit to the bones). The technical approach of the system is to use computed tomography (CT) for three-dimensional (3-D) planning and a robot for automated bone milling. The planning (surgical CAD) is performed on the ORTHODOC workstation, which enables the surgeon to graphically position a 3-D model of the prosthesis (or prostheses) with respect to the CT image, thereby creating a surgical plan. In the operating room (Surgical CAM), the robot is registered to the CT image so that the surgical plan can be transformed from the CT coordinate system to the robot coordinate system. The robot then machines the bone according to the plan, using a high-speed milling tool.

Surgical Assistance

Medical interventions are highly interactive processes, and many critical decisions are made in the operating room and




Figure 1. Architecture of a surgical CAD/CAM system, where the preoperative phase is CAD, the intraoperative phase is CAM, and the postoperative phase is TQM.




Figure 2. The ROBODOC system for orthopedic surgery. (a) The robot is being used for total hip replacement surgery. (b) Close-up of robotic milling of femur.

The da Vinci telerobotic system has recently become more widely employed for minimally invasive surgery.

executed immediately. The goal of computer-assisted medical systems, including surgical robots, is not to replace the physician with a machine but, rather, to provide intelligent, versatile tools that augment the physician's ability to treat patients. There are many forms of technological assistance. In this section, we focus on robotic assistance. Some nonrobotic technologies are reviewed in the "Other Technologies" section.

There are two basic augmentation strategies: 1) improving the physician's existing sensing and/or manipulation, and 2) increasing the number of sensors and manipulators available to the physician (e.g., more eyes and hands). In the first case, the system can give even average physicians superhuman capabilities such as X-ray vision, elimination of hand tremor, or the ability to perform dexterous operations inside the patient's body. A special subclass is remote telesurgery systems, which permit the physician to operate on patients at distances ranging from a few meters to several thousand kilometers. In the second case, the robot operates side by side with the physician and performs functions such as endoscope holding, tissue retraction, or limb positioning. These systems typically provide one or more direct control interfaces such as joysticks, head trackers, or voice control but could also include intelligence to demand less of the physician's attention during use.

The da Vinci system (Intuitive Surgical, Inc.) is a telesurgery system that demonstrates both of these augmentation approaches [2]. As shown in Figure 3, the system consists of a patient-side slave robot and a master control console. The slave robot has three or four robotic arms that manipulate a stereo endoscope and dexterous surgical instruments such as scissors, grippers, and needle holders. The surgeon sits at the master control console and grasps handles attached to two dexterous master manipulator arms, which are capable of exerting limited

amounts of force feedback to the surgeon. The surgeon's hand motions are sensed by the master manipulators, and these motions are replicated by the slave manipulators. A variety of control modes may be selected via foot pedals on the master console and used for such purposes as determining which slave arms are associated with the hand controllers. Stereo video is transmitted from the endoscope to a pair of high-quality video monitors in the master control console, thus providing high-fidelity stereo visualization of the surgical site. The display and master manipulators are arranged so that it appears to the surgeon that the surgical instruments (inside the patient) are in the same position as his or her hands inside the master control console. Thus, the da Vinci system improves the surgeon's eyes and hands by enabling them to (remotely) see and manipulate tissue inside the patient through incisions that are too small for direct visualization and manipulation. By providing three or four slave robot arms, the da Vinci system also endows the surgeon with more than two hands.

Other Technologies

Robotics is not the only manner in which computers can be used to assist medical procedures. One important, and widely used, alternative is a navigation system, which consists of a sensor (tracker) that can measure the position and orientation of instruments in 3-D space (typically, the instruments contain special tracker targets). If the tracker coordinate system is registered to a preoperative or intraoperative image (see the "Registration" section), the navigation system can display the position and orientation of the instrument with respect to the image. This improves the physician's visualization by enabling him or her to see the internal structure, molecular information, and/or functional data, depending on the type of image. This can also enable the physician to execute a preoperative plan (surgical CAD/CAM), e.g., by aligning an instrument with respect to a target defined in the preoperative image. Currently, the most widely used tracking technology is optical because of its relatively high accuracy, predictable performance, and insensitivity to environmental variations. The primary limitation of optical trackers is that they require a clear line of sight between the camera and the instruments being tracked. This precludes their use for instruments inside the body. Electromagnetic tracking systems are free from line-of-sight constraints but are generally less accurate, especially due to field distortions caused by metallic objects.

Technology and Challenges

Surgical robots present a unique set of design challenges due to the requirements for miniaturization, safety, sterility, and adaptation to changing conditions. This section reviews current practices and challenges in manipulation, sensing, registration, user interfaces, and system design.

Manipulation

Surgical robots must satisfy requirements not found in industrial robotics. They must operate safely in a workspace shared with humans; they generally must operate in a sterile environment; and they often require high dexterity in small spaces. An




Figure 3. The da Vinci surgical system (courtesy Intuitive Surgical, Inc.).

additional challenge occurs when the robot must operate in the proximity of a magnetic resonance imaging (MRI) scanner, whose high magnetic field precludes the use of many conventional robotic components.

The topic of safety design is covered in detail in the “Safety Design” section. There are, however, certain safety factors that should be considered during the design of a surgical manipulator. First, unlike industrial robots, where speed and strength are desirable attributes, a surgical robot should only be as fast and strong as needed for its intended use. In most cases, the robot should not be capable of moving faster or with more force than the physician. An obvious exception could occur for a robot that operates on a rapidly moving organ, such as a beating heart. Even in this case, there are innovative solutions that do not require rapid motion, such as Heartlander [6], which is designed to attach to a beating heart using suction and move along it with inchworm locomotion. Another safety-related design parameter is the robot’s workspace, which ideally should only be as large as needed. This is difficult to achieve in practice, given the high variability between patients and the differences in the way that physicians perform procedures. Some researchers have reported parallel manipulators, which have smaller workspaces (and higher rigidity) than serial robots [7]–[10].

Sterility is a major design challenge. It is not easy to design reusable devices that can withstand multiple sterilization cycles. One common solution is to create a disposable device that only needs to be sterilized once, usually by the manufacturer. This is practical for low-cost parts. Another issue with a reusable device is that it must be cleaned between procedures. Thus, crevices that can trap blood or other debris should be avoided. The most common approach is to design the surgical robot so that its end effector (or tool) can be removed and sterilized, while the rest of the robot is covered with a disposable sterile drape or bag (e.g., as illustrated for ROBODOC in Figure 2). This is particularly difficult when the end effector or tool includes electromechanical components.

Size matters for surgical robots. Operating rooms and interventional suites are usually small, and, thus, a large robot can take too much space. This has been a complaint for many commercially available systems, such as daVinci and ROBODOC, which are large floor-standing robots. In orthopedics, there have been recent examples of smaller, bone-mounted robots [7]–[9]. Size is especially critical when the robot, or part of it, must work inside the body. For example, although the da Vinci system is large, its robotic EndoWrist tools, with diameters from 5–8 mm, are a marvel of miniaturization and can pass into the body via small entry ports.

The design of MRI-compatible robots is especially challenging because MRI relies on a strong magnetic field and radio frequency (RF) pulses, and so it is not possible to use components that can interfere with, or be susceptible to, these physical effects. This rules out most components used for robots, such as electric motors and ferromagnetic materials. Thus, MRI-compatible robots typically use nonmetallic links and piezoelectric, pneumatic, or hydraulic motors. This topic will be discussed in greater detail in a subsequent part of this tutorial.

Sensing

Besides internal sensors, such as joint encoders, a surgical robot often needs external sensors to enable it to adapt to its relatively unstructured and changing environment. Common examples are force sensors and vision systems, which translate naturally into the human senses of touch and sight. For this reason, they are often used for surgical assistants. For example, the da Vinci system provides exquisite stereo video feedback, although it is often criticized for not providing force feedback (a component of haptic feedback). Without force feedback, the surgeon must use visual cues, such as the tautness of a suture or the deflection of tissue, to estimate the forces. If these cues are misread, the likely outcome is a broken suture or damaged tissue [11].

Real-time imaging such as ultrasound, spectroscopy, and optical coherence tomography (OCT) can provide significant benefits when they enable the physician to see subsurface structures and/or tissue properties. For example, when resecting a brain tumor, this type of sensing can alert the surgeon before he or she accidentally cuts a major vessel that is obscured by the tumor. Preoperative images, when registered to the robot, can potentially provide this information, but only if the anatomy does not change significantly during the procedure. This is rarely the case, except when working with rigid structures such as bones. Once again, it is necessary to overcome challenges in sterility and miniaturization to provide this sensing where it is needed, which is usually at or near the instrument tip.

Sensors that directly measure physiologic properties, such as tissue oxygenation, are also useful. For example, a smart retractor that uses pulse oxymetry principles to measure the oxygenation of blood can detect the onset of ischemia (insufficient blood flow) before it causes a clinical complication [12].

Registration

Geometric relationships between portions of the patient’s anatomy, images, robots, sensors, and equipment are fundamental to all areas of computer-integrated medicine. There is an extensive literature on techniques for determining the transformations between the associated coordinate systems [13], [14]. Given two coordinates $\vec{v}_A = [x_A, y_A, z_A]$ and $\vec{v}_B = [x_B, y_B, z_B]$ corresponding to comparable features in two coordinate systems Ref_A and Ref_B , the process of registration is simply that of finding a function $\mathbf{T}_{AB}(\cdots)$ such that

$$\vec{v}_B = \mathbf{T}_{AB}(\vec{v}_A).$$

Although nonrigid registrations are becoming more common, $\mathbf{T}_{AB}(\cdots)$ is still usually a rigid body transformation of the form

$$\vec{v}_B = \mathbf{T}_{AB}(\vec{v}_A) = \mathbf{R}_{AB} \bullet \vec{v}_A + \vec{p}_{AB},$$

where $\mathbf{R}_{AB}$ represents a rotation and $\vec{p}_{AB}$ represents a translation.

The simplest registration method is a paired-point registration in which a set of N points ($N \geq 3$) is found in the first coordinate system and matched (one to one) with N

The development of surgical robots is motivated primarily by the desire to enhance the effectiveness of a procedure.

corresponding points in the second coordinate system. The problem of finding the transformation that best aligns the two sets of points is often called the Procrustes problem, and there are well-known solutions based on quaternions [15] and rotation matrices [16], [17]. This method works best when it is possible to identify distinct points in the image and on the patient. This is usually straightforward when artificial fiducials are used. For example, ROBODOC initially used a fiducial-based registration method, with three metal pins (screws) inserted into the bone prior to the CT scan. It was easy to locate the pins in the CT image, via image processing, due to the high contrast between metal and bone. Similarly, it was straightforward for the surgeon to guide the robot's measurement probe to physically contact each of the pins.

Point-to-surface registration methods can be employed when paired-point registration is not feasible. Typically, this involves matching a cloud of points that is collected intraoperatively to a 3-D surface model that is constructed from the pre-operative image. The most widely used method is iterative closest point (ICP) [18]. Briefly, ICP starts with an initial guess of the transformation, which is used to transform the points to the same coordinate system as the surface model. The closest points on the surface model are identified and a paired-point registration method is used to compute a new estimate of the transformation. The process is repeated with the new transformation until a termination condition is reached. Although ICP often works well, it is sensitive to the initial guess and can fail to find the best solution if the guess is poor. Several ICP variations have been proposed to improve its robustness in this case, and other techniques, such as an unscented Kalman filter [19], have recently been proposed. These methods can also be used for surface-to-surface registration by sampling one of the surfaces.

Nonrigid (elastic or deformable) registration is often necessary because many parts of the anatomy (e.g., soft tissue and organs) change shape during the procedure. This is more difficult than rigid registration and remains an active area of research. To date, most surgical CAD/CAM systems have been applied to areas such as orthopedics, where deformations are small and rigid registration methods can be employed.

User Interfaces and Visualization

Standard computer input devices, such as keyboards and mice, are generally inappropriate for surgical or interventional environments because it is difficult to use them in conjunction with other medical instrumentation and maintain sterility. Foot pedals are often used because they do not interfere with whatever the physician is doing with his or her hands, and they

do not require sterilization. Handheld pendants (button boxes) are also used; in this case, the pendant is either sterilized or covered by a sterile drape. It is important to note, however, that the robot itself can often provide a significant part of the user interface. For example, the da Vinci system relies on the two master manipulators (one for each hand), with foot pedals to change modes. The ROBODOC system not only includes a five-button pendant to navigate menus but also uses a force-control (hand guiding) mode that enables the surgeon to manually move the robot.

Computer output is traditionally provided by graphical displays. Fortunately, these can be located outside the sterile field. Unfortunately, the ergonomics are often poor because the physician must look away from the operative site (where his or her hands are manipulating the instruments) to see the computer display. Some proposed solutions include heads-up displays, image overlay systems [20], [21], and lasers, which project information onto the operative field [22].

Surgical Robot System Design

A surgical robot includes many components, and it is difficult to design one from scratch. There is no off-the-shelf surgical robot for research, and it is unlikely that one robot or family of robots will ever satisfy the requirements of the diaspora of medical procedures. In the software realm, however, there are open source software packages that can help. The most mature packages are for medical image visualization and processing, particularly the Visualization Toolkit (VTK, www.vtk.org) and the Insight Toolkit (ITK, www.itk.org). Customizable applications, such as 3-D Slicer (www.slicer.org), package VTK, ITK, and a plethora of research modules.

Few packages exist for computer-assisted interventions. The Image Guided Surgery Toolkit (IGSTK, www.igstk.org) enables researchers to create a navigation system by connecting a tracking system to a computer. At Johns Hopkins University, we are creating a software framework for a surgical assistant workstation (SAW), based on our Computer-Integrated Surgical Systems and Technology (CISST) libraries [23] (www.cisst.org), which focus on the integration of robot control and real-time sensing with the image processing and visualization toolkits described previously.

Surgical Robot Design Process

This section presents a detailed discussion of the risk analysis, safety design, and validation phases of the design process. Although these topics are not unique to surgical robots, they are obviously of extreme importance.

Risk Analysis

Safety is an important consideration for both industrial and surgical robots [24]. In an industrial setting, safety can often be achieved by keeping people out of the robot's workspace or by shutting down the system if a person comes too close. In contrast, for surgical robots it is generally necessary for human beings, including the patient and the medical staff, to be inside the robot's workspace. Furthermore, the robot may be holding a potentially dangerous device, such as a cutting instrument,

that is supposed to actually contact the patient (in the correct place, of course). If the patient is anesthetized, it is not possible for him or her to actively avoid injury.

Proper safety design begins with a risk (or hazard) analysis. A failure modes effects analysis (FMEA) or failure modes effects and criticality analysis (FMECA) are the most common methods [25]. These are bottom-up analyses, where potential component failures are identified and traced to determine their effect on the system. Methods of control are devised to mitigate the hazards associated with these failures. The information is generally presented in a tabular format (see Table 1). The FMECA adds the criticality assessment, which consists of three numerical parameters: the severity (S), occurrence (O), and detectability (D) of the failure. A risk priority number (RPN) is computed from the product of these parameters, which determines whether additional methods of control are required. The FMEA/FMECA is a proactive analysis that should begin early in the design phase and evolve as hazards are identified and methods of control are developed. Another popular method is a fault tree analysis (FTA), which is a top-down analysis and is generally more appropriate for analyzing a system failure after the fact.

Safety Design

As an illustrative example of how to apply these methods in the design phase, consider a multilink robot system where each link is driven by a feedback-controlled motor, as shown in Figure 4. The error, $e(t)$, between the desired position $x_d(t)$ and the measured position $x_a(t)$ is computed and used to determine the control output $u(t)$ that drives the motor. An encoder failure will cause the system to measure a persistent steady-state error and therefore continue to drive the motor to attempt to reduce this error. An amplifier failure can cause it to apply an arbitrary voltage to the motor that is independent of the control signal $u(t)$. The controller will sense the increasing error and adjust $u(t)$ to attempt to compensate, but this will have no effect.

These failure modes are shown in the FMEA presented in Table 1. The result in both these cases is that the robot will move until it hits something (typically, the effect on system is more descriptive and includes application-specific information, such as the potential harm to the patient). This is clearly unacceptable for a surgical robot, and so methods of control are necessary. One obvious solution, shown in Table 1, is to allow the control software to disable the motor power, via a relay, whenever the error, $e(t)$, exceeds a specified threshold. This will prevent a catastrophic, headline-grabbing runaway robot scenario, but is the robot safe enough for surgical use? The answer is that it depends on the application and on the physical parameters of the system. To illustrate this, consider the case where the power amplifier fails and applies maximum voltage to the motor. As shown in Figure 5, if E is the error threshold (i.e., the point at which the control software disables motor power via the relay), the final joint position error, $\Delta P_{\max}$, is given by

$E + V_{\max} \times \Delta T + \Delta P_{\text{off}}$, where ΔT is the control period, $V_{\max}$ is the maximum joint velocity (assuming the robot had sufficient time to accelerate), and ΔP_{off} is the distance the robot travels after power off due to inertia or external forces. The actual value of $\Delta P_{\max}$ depends on the robot design, but it is not uncommon for this to be several millimeters. Although a one-time glitch of this magnitude may be tolerable for some surgical procedures, it is clearly not acceptable in others. In those cases, it is necessary to make design modifications to decrease $\Delta P_{\max}$, e.g., by decreasing $V_{\max}$, or to forgo the use of an active robot. This safety analysis was a prime motivation for researchers who developed passive robots such as Cobots [26] and PADyC [27].

There are safety issues that must be considered regardless of whether a robot is active or passive. One example occurs when the robot's task is to accurately position an instrument




Figure 4. Computer control of a robot joint, showing the motor (M), encoder (E), and power amplifier (Amp).




Figure 5. Illustration of maximum possible error: E is the error threshold, $V_{\max}$ is the maximum velocity, ΔT is the control period, and ΔP_{off} is the robot stopping distance.

Table 1. Excerpt from a sample FMEA.			
Failure Mode	Effect on System	Causes	Methods of Control*
Robot out of control	Robot may hit something	Encoder failure, broken wire	Trip relay when error tolerance exceeded
Robot out of control	Robot may hit something	Amplifier failure	Trip relay when error tolerance exceeded
*Methods of control can initially be empty and then populated during the design phase.			

Surgical robots present a unique set of design challenges due to the requirements for miniaturization, safety, sterility, and adaptation to changing conditions.

or instrument guide. The position of a robot-held tool is typically determined by applying the robot's forward kinematic equations to the measured joint positions. An inaccurate joint sensor (e.g., an incremental encoder that intermittently gains or loses counts) can cause a large position error. One method of control is to introduce a redundant sensor and use software to verify whether both sensors agree within a specified tolerance. Practical considerations dictate the need for a tolerance to account for factors such as mechanical compliance between the sensors and differences in sensor resolution and time of data acquisition. This limits the degree with which accuracy can be assured. Note also that although redundant sensors remove one single point of failure (i.e., sensor failure), it is necessary to avoid a single point of failure in the implementation. For example, if both sensors are placed on the motor shaft, they cannot account for errors in the joint transmission, e.g., due to a slipped belt.

A final point is that redundancy is not sufficient if failure of one component cannot be detected. For example, consider the case where the robot is holding a pneumatic cutting tool, and a solenoid is used to turn the tool on and off. If the solenoid fails in the open (on) state, the cutting tool may be activated at an unsafe time. It is tempting to address this hazard by putting a second solenoid in series with the first, as shown in Figure 6. This is not an acceptable solution, however, because if one solenoid fails in the open state, the system will appear to operate correctly (i.e., the software can still turn the cutter on or off). Therefore, this system once again has a single point of failure. This is not a hypothetical scenario—it actually appeared in the risk analysis for the ROBODOC system, which uses a pneumatic cutting tool. The concern was that a failed solenoid could cause

injury to the surgeon if the failure occurred while the surgeon was inserting or removing the cutting bit. ROBODOC adopted a simple method of control, which was to display a screen instructing the surgeon to physically disconnect the pneumatic supply prior to any cutting tool change.

Validation

Validation of computer-integrated systems is challenging because the key measure is how well the system performs in an operating room or interventional suite with a real patient. Clearly, for both ethical and regulatory reasons, it is not possible to defer validation until a system is used with patients. Furthermore, it is difficult to quantify intraoperative performance because there are limited opportunities for accurate postoperative assessment. For example, CT scans may not provide sufficient contrast for measuring the postoperative result, and they expose the patient to additional radiation. For these reasons, most computer-integrated systems are validated using phantoms, which are objects that are designed to mimic (often very crudely) the relevant features of the patient.

One of the key drivers of surgical CAD/CAM is the higher level of accuracy that can be achieved using some combination of computers, sensors, and robots. Therefore, it is critical to evaluate the overall accuracy of such a system. One common technique is to create a phantom with a number of features (e.g., fiducials) whose locations are accurately known, either by precise manufacturing or measurement. Some of these features should be used for registration, whereas others should correspond to targets. The basic technique is to image the phantom, perform the registration, and then locate the target features. By convention, the following types of error are defined [28] as follows:

- ♦ *fiducial localization error (FLE)*: the error in locating a fiducial in a particular coordinate system (i.e., imaging system or robot system)
- ♦ *fiducial registration error (FRE)*: the root mean square (RMS) residual error at the registration fiducials, i.e.,

$$FRE = \sqrt{\frac{1}{N} \sum_{k=1}^N \|\vec{\mathbf{b}}_k - \mathbf{T} \bullet \vec{\mathbf{a}}_k\|^2}$$

where $\mathbf{T}$ is the registration transform and $(\vec{\mathbf{a}}_k, \vec{\mathbf{b}}_k)$ are matched pairs of homologous fiducials ($k = 1, \dots, N$)

- ♦ *target registration error (TRE)*: the error in locating a feature or fiducial that was not used for the registration; if multiple targets are available, the mean error is often reported as the TRE.

Although it is necessary to validate that a surgical robot meets its requirements, including those related to accuracy, it is important to realize that higher accuracy may not lead to a clinical benefit. Validation of clinical utility is often possible only via clinical trials.

Summary

This article presents the first of a three-part tutorial on surgical and interventional robotics. The core concept is that a surgical robot couples information to action in the operating room or




Figure 6. Example of poorly designed redundant system. The second solenoid does not provide sufficient safety because the system cannot detect when either solenoid has failed in the open state.

interventional suite. This leads to several potential benefits, including increased accuracy and the ability to intervene in areas that are not accessible with conventional instrumentation. We defined the categories of surgical CAD/CAM and surgical assistance. The former is intended to accurately execute a defined plan. The latter is focused on providing augmented capabilities to the physician, such as superhuman or auxiliary (additional) eyes and hands. These categories will be the focus of the final two parts of this tutorial.

There are numerous challenges in surgical manipulation, sensing, registration, user interfaces, and system design. Many of these challenges result from the requirements for safety, sterility, small size, and adaptation to a relatively unstructured (and changing) environment. Some software toolkits are available to facilitate the design of surgical robotics systems.

The design of a surgical robot should include a risk analysis. Established methodologies such as FMEA/FMECA can be used to identify potential hazards. Safety design should consider and eliminate single points of failure whenever possible. Validation of system performance is critical but is complicated by the difficulty of simulating realistic clinical conditions.

Surgical robotics is a challenging field, but it is rewarding because the ultimate goal is to improve the health and quality of human life.

Acknowledgments

The authors gratefully acknowledge the National Science Foundation for supporting our work in this field through the Engineering Research Center for Computer-Integrated Surgical Systems and Technology (CISST ERC), NSF Grant EEC 9731748. Related projects have also been supported by Johns Hopkins University, the National Institutes of Health, the Whitaker Foundation, the Department of Defense, and our CISST ERC industrial affiliates.

Keywords

Surgical robots, medical robots, robot safety.

References

- [1] Y. S. Kwok, J. Hou, E. A. Jonckheere, and S. Hayati, "A robot with improved absolute positioning accuracy for CT-guided stereotactic brain surgery," *IEEE Trans. Biomed. Eng.*, vol. 35, no. 2, pp. 153–160, 1988.
- [2] G. S. Guthart and J. K. Salisbury, "The intuitive telesurgery system: Overview and application," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA 2000)*, San Francisco, vol. 1, pp. 618–621.
- [3] R. H. Taylor and D. Stoianovici, "Medical robotics in computer-integrated surgery," *IEEE Trans. Robot. Automat.*, vol. 19, no. 3, pp. 765–781, 2003.
- [4] R. H. Taylor, B. D. Mittelstadt, H. A. Paul, W. Hanson, P. Kazanzides, J. F. Zuhars, B. Williamson, B. L. Musits, E. Glassman, and W. L. Bargar, "An image-directed robotic system for precise orthopaedic surgery," *IEEE Trans. Robot. Automat.*, vol. 10, no. 3, Jun. 1994.
- [5] P. Kazanzides, "Robot assisted surgery: The ROBODOC experience," in *Proc. 30th Int. Symp. Robotics (ISR)*, Tokyo, Japan, Nov. 1999, pp. 261–286.
- [6] N. Patronik, C. Riviere, S. E. Qarra, and M. A. Zenati, "The Heart-Lander: A novel epicardial crawling robot for myocardial injections," in *Proc. 19th Int. Congr. Computer Assisted Radiology and Surgery*, 2005, vol. 1281C, pp. 735–739.

- [7] M. Shoham, M. Burman, E. Zehavi, L. Joskowicz, E. Batkalin, and Y. Kunicher, "Bone-mounted miniature robot for surgical procedures: Concept and clinical applications," *IEEE Trans. Robot. Automat.*, vol. 19, no. 5, pp. 893–901, Oct. 2003.
- [8] A. Wolf, B. Jaramaz, B. Lisien, and A. M. DiGioia, "MBARS: Mini bone-attached robotic system for joint arthroplasty," *Int. J. Med. Robot. Comp. Assist. Surg.*, vol. 1, no. 2, pp. 101–121, Jan. 2005.
- [9] J. H. Chung, S. Y. Ko, D. S. Kwon, J. J. Lee, Y. S. Yoon, and C. H. Won, "Robot-assisted femoral stem implantation using an intramedulla gauge," *IEEE Trans. Robot. Automat.*, vol. 19, no. 5, pp. 885–892, Oct. 2003.
- [10] G. Brandt, A. Zimolong, L. Carrat, P. Merloz, H. W. Staudte, S. Lavalée, K. Radermacher, and G. Rau, "CRIGOS: A compact robot for image-guided orthopedic surgery," *IEEE Trans. Inform. Technol. Biomed.*, vol. 3, no. 4, pp. 252–260, Dec. 1999.
- [11] A. M. Okamura, "Methods for haptic feedback in teleoperated robot-assisted surgery," *Ind. Robot.*, vol. 31, no. 6, pp. 499–508, 2004.
- [12] G. Fischer, T. Akinbiyi, S. Saha, J. Zand, M. Talamini, M. Marohn, and R. H. Taylor, "Ischemia and force sensing surgical instruments for augmenting available surgeon information," in *Proc. IEEE Int. Conf. Biomedical Robotics and Biomechanics (BioRob 2006)*, Pisa, Italy, 2006, pp. 1030–1035.
- [13] J. B. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, no. 1, pp. 1–37, 1998.
- [14] S. Lavalée, "Registration for computer-integrated surgery: methodology, state of the art," in *Computer-Integrated Surgery*, R. H. Taylor, S. Lavalée, G. Burdea, and R. Mosges, Eds. Cambridge, MA: MIT Press, 1996, pp. 77–98.
- [15] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Am. A*, vol. 4, no. 4, pp. 629–642, 1987.
- [16] K. Arun, T. Huang, and S. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 9, no. 5, pp. 698–700, 1987.
- [17] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 4, pp. 376–380, 1991.
- [18] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [19] M. H. Moghari and P. Abolmaesumi, "Point-based rigid-body registration using an unscented Kalman filter," *IEEE Trans. Med. Imag.*, vol. 26, no. 12, pp. 1708–1728, Dec. 2007.
- [20] M. Blackwell, C. Nikou, A. M. DiGioia, and T. Kanade, "An image overlay system for medical data visualization," *Med. Image Anal.*, vol. 4, no. 1, pp. 67–72, 2000.
- [21] G. Fichtinger, A. Deguet, K. Masamune, E. Balogh, G. S. Fischer, H. Mathieu, R. H. Taylor, S. J. Zinreich, and L. M. Fayad, "Image overlay guidance for needle insertion on CT scanner," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 8, pp. 1415–1424, Aug. 2005.
- [22] T. Sasama, N. Sugano, Y. Sato, Y. Momoi, T. Koyama, Y. Nakajima, I. Sakuma, M. G. Fujie, K. Yonenobu, T. Ochi, and S. Tamura, "A novel laser guidance system for alignment of linear surgical tools: Its principles and performance evaluation as a man-machine system," in *Proc. 5th Int. Conf. Medical Image Computing and Computer-Assisted Intervention*, 2002, vol. 2489, pp. 125–132.
- [23] A. Kapoor, A. Deguet, and P. Kazanzides, "Software components and frameworks for medical robot control," in *Proc. IEEE Conf. Robotics and Automation (ICRA)*, Orlando, FL, May 2006, pp. 3813–3818.
- [24] B. Davies, "A discussion of safety issues for medical robots," in *Computer-Integrated Surgery*, R. Taylor, S. Lavalée, G. Burdea, and R. Moesges, Eds. Cambridge, MA: MIT Press, 1996, pp. 287–296.
- [25] R. E. McDermott, R. J. Mikulak, and M. R. Beauregard, "The Basics of FMEA," New York, Quality Resources, 1996.

- [26] M. A. Peshkin, J. E. Colgate, W. Wannasuphoprasit, C. A. Moore, R. B. Gillespie, and P. Akella, "Cobot architecture," *IEEE Trans. Robot. Automat.*, vol. 17, no. 4, pp. 377–390, Aug. 2001.
- [27] O. Schneider and J. Troccaz, "A six-degree-of-freedom passive arm with dynamic constraints (PADyC) for cardiac surgery applications: Preliminary experiments," *Comput. Aided Surg.*, vol. 6, no. 6, pp. 340–351, 2001.
- [28] C. Maurer, J. Fitzpatrick, M. Wang, R. Galloway, R. Maciunas, and G. Allen, "Registration of head volume images using implantable fiducial markers," *IEEE Trans. Med. Imag.*, vol. 16, no. 4, pp. 447–462, Aug. 1997.

Peter Kazanzides received the B.S., M.S., and Ph.D. degrees in electrical engineering from Brown University in 1983, 1985, and 1988, respectively. He worked on surgical robotics in March 1989 as a postdoctoral researcher at the International Business Machines (IBM) T.J. Watson Research Center. He cofounded Integrated Surgical Systems (ISS) in November 1990 to commercialize the robotic hip replacement research performed at IBM and the University of California, Davis. As the director of robotics and software, he was responsible for the design, implementation, validation and support of the ROBODOC System. He joined the Engineering Research Center for Computer-Integrated Surgical Systems and Technology (CISST ERC) in December 2002, and currently, he is an assistant research professor of computer science at Johns Hopkins University.

Gabor Fichtinger received his B.S. and M.S. degrees in electrical engineering and his Ph.D. degree in computer science from the Technical University of Budapest, Hungary, in 1986, 1988, and 1990, respectively. He has developed image-guided surgical interventional systems. He specializes in robot-assisted image-guided needle-placement procedures, primarily for cancer diagnosis and therapy. He is an associate professor of computer science, electrical engineering, mechanical engineering, and surgery at Queen's University, Canada, with adjunct appointments at the Johns Hopkins University.

Gregory D. Hager is a professor of computer science at Johns Hopkins University. He received the B.A. degree, *summa cum laude*, in computer science and mathematics from Luther College, in 1983, and the M.S. and Ph.D. degrees in computer science from the University of Pennsylvania in 1985 and 1988, respectively. From 1988 to 1990, he was a Fulbright junior research fellow at the University of Karlsruhe and the Fraunhofer Institute IITB in Karlsruhe, Germany. From 1991 to 1999, he was with the Computer Science Department at Yale University. In 1999, he joined the Computer Science Department at Johns Hopkins University, where he is the deputy director of the Center for Computer Integrated Surgical Systems and Technology. He has authored more than 180 research articles and books in the area of robotics and computer vision. His current

research interests include visual tracking, vision-based control, medical robotics, and human-computer interaction. He is a Fellow of the IEEE.

Allison M. Okamura received the B.S. degree from the University of California at Berkeley, in 1994, and the M.S. and Ph.D. degrees from Stanford University in 1996 and 2000, respectively, all in mechanical engineering. She is currently an associate professor of mechanical engineering and the Decker Faculty Scholar at Johns Hopkins University. She is the associate director of the Laboratory for Computational Sensing and Robotics and a thrust leader of the National Science Foundation Engineering Research Center for Computer-Integrated Surgical Systems and Technology. Her awards include the 2005 IEEE Robotics Automation Society Early Academic Career Award, the 2004 National Science Foundation Career Award, the 2004 Johns Hopkins University George E. Owen Teaching Award, and the 2003 Johns Hopkins University Diversity Recognition Award. Her research interests include haptics, teleoperation, medical robotics, virtual environments and simulators, prosthetics, rehabilitation engineering, and engineering education.

Louis L. Whitcomb completed his B.S. and Ph.D. degrees at Yale University in 1984 and 1992, respectively. His research focuses on the design, dynamics, navigation, and control of robot systems. He has numerous patents in the field of robotics, and he is a Senior Member of the IEEE. He is the founding director of the Johns Hopkins University Laboratory for Computational Sensing and Robotics. He is a professor at the Department of Mechanical Engineering, with joint appointment in the Department of Computer Science, at the Johns Hopkins University.

Russell H. Taylor received his Ph.D. degree in computer science from Stanford in 1976. He joined IBM Research in 1976, where he developed the AML robot language and managed the Automation Technology Department and (later) the Computer-Assisted Surgery Group before moving in 1995 to Johns Hopkins University, where he is a professor of computer science, with joint appointments in mechanical engineering, radiology and surgery. He is the Director of the NSF Engineering Research Center for Computer-Integrated Surgical Systems and Technology. He is the author of more than 200 refereed publications. He is a Fellow of the IEEE and AIMB and is a recipient of the Maurice Müller award for excellence in computer-assisted orthopedic surgery.

Address for Correspondence: Peter Kazanzides, Department of Computer Science, CSEB 120, Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA. E-mail: pkaz@jhu.edu.